The 16kB Single-Cycle Read-Access Cache on a Next-Generation 64b Itanium™ Microprocessor

D. Bradley, P. Mahoney, B. Stackhouse
Hewlett Packard, Fort Collins, CO
Intel Corp., Fort Collins, CO

Presentation Overview

- Cache Overview
- Prevalidated Tag Microarchitecture
- Three Data Arrays
  - Level 1 TLB Physical Architecture
  - Prevalidated Tag Array Physical Architecture
  - Data Cache Design
- Performance
- Q & A

Cache Overview

- 16 kB of data (4 ways × 4 kB/way)
  - 16 times the size (in bits) of the integer register file
- Prevalidated tags for fast loads
- 4 Ports (2 load ports + 2 store ports)
  - Load ports are independent.
  - Each store port has a 1 in 8 chance of conflicting with each other valid port.
- “True” single cycle load access, including:
  - address translation
  - data delivery
  - data read
  - integer unit data bypass
- Physically addressed (50 physical address bits)
- Write-through to Level 2 Cache
- 4kB way size (64B line × 64 indices)

Prevalidated Tags

- Traditional physically-addressed caches have the following serialized timing path:

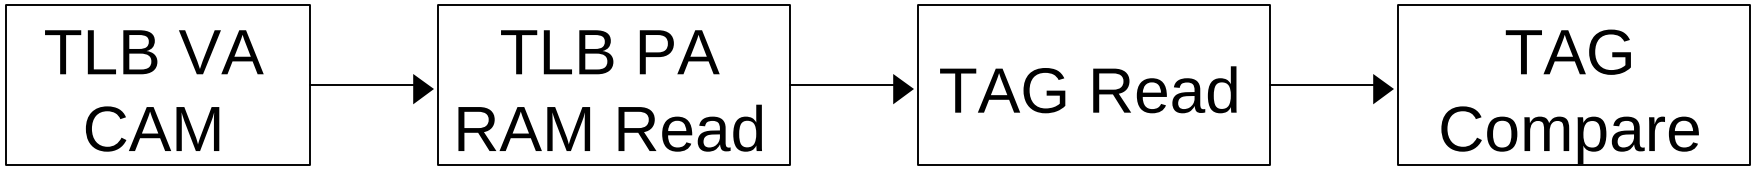

- On McKinley, L1 way size <= page size. Thus, the index bits into the tag array and data array are **untranslated**.
- Instead of storing physical address-based tags, McKinley stores a 32b one-hot pointer to a TLB entry to stand-in for the PA already contained there. This is what is meant by "**prevalidation**".
- Example:
  - Tlb entry 5: VA:0xa5a5a5a5a5a5axxx ->PA:0x0b4b4b4b4bxxx
  - Conventional tag: 0x0b4b4b4b4b (38 bits)
  - Prevalidated one-hot tag: 0x00000020 (32 bits)

Prevalidated Tag Advantages

- Faster TLB outputs (no TLB RAM read)

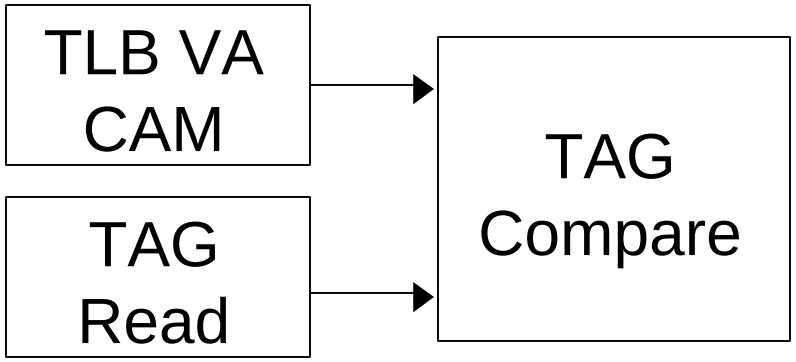

- Faster TAG compare
  - Dual rail compare + 38-bit dynamic AND vs.
  - Dual rail miscompare + 38-bit OR + a hard clock edge vs.
  - **Single rail compare + 32-bit dynamic OR**
- Less wiring in both the TLB and TAG arrays
  - 38 address bits × 2 rails per bit vs.
  - **32 bits × 1 rail per bit**

Block Diagram (1 port)

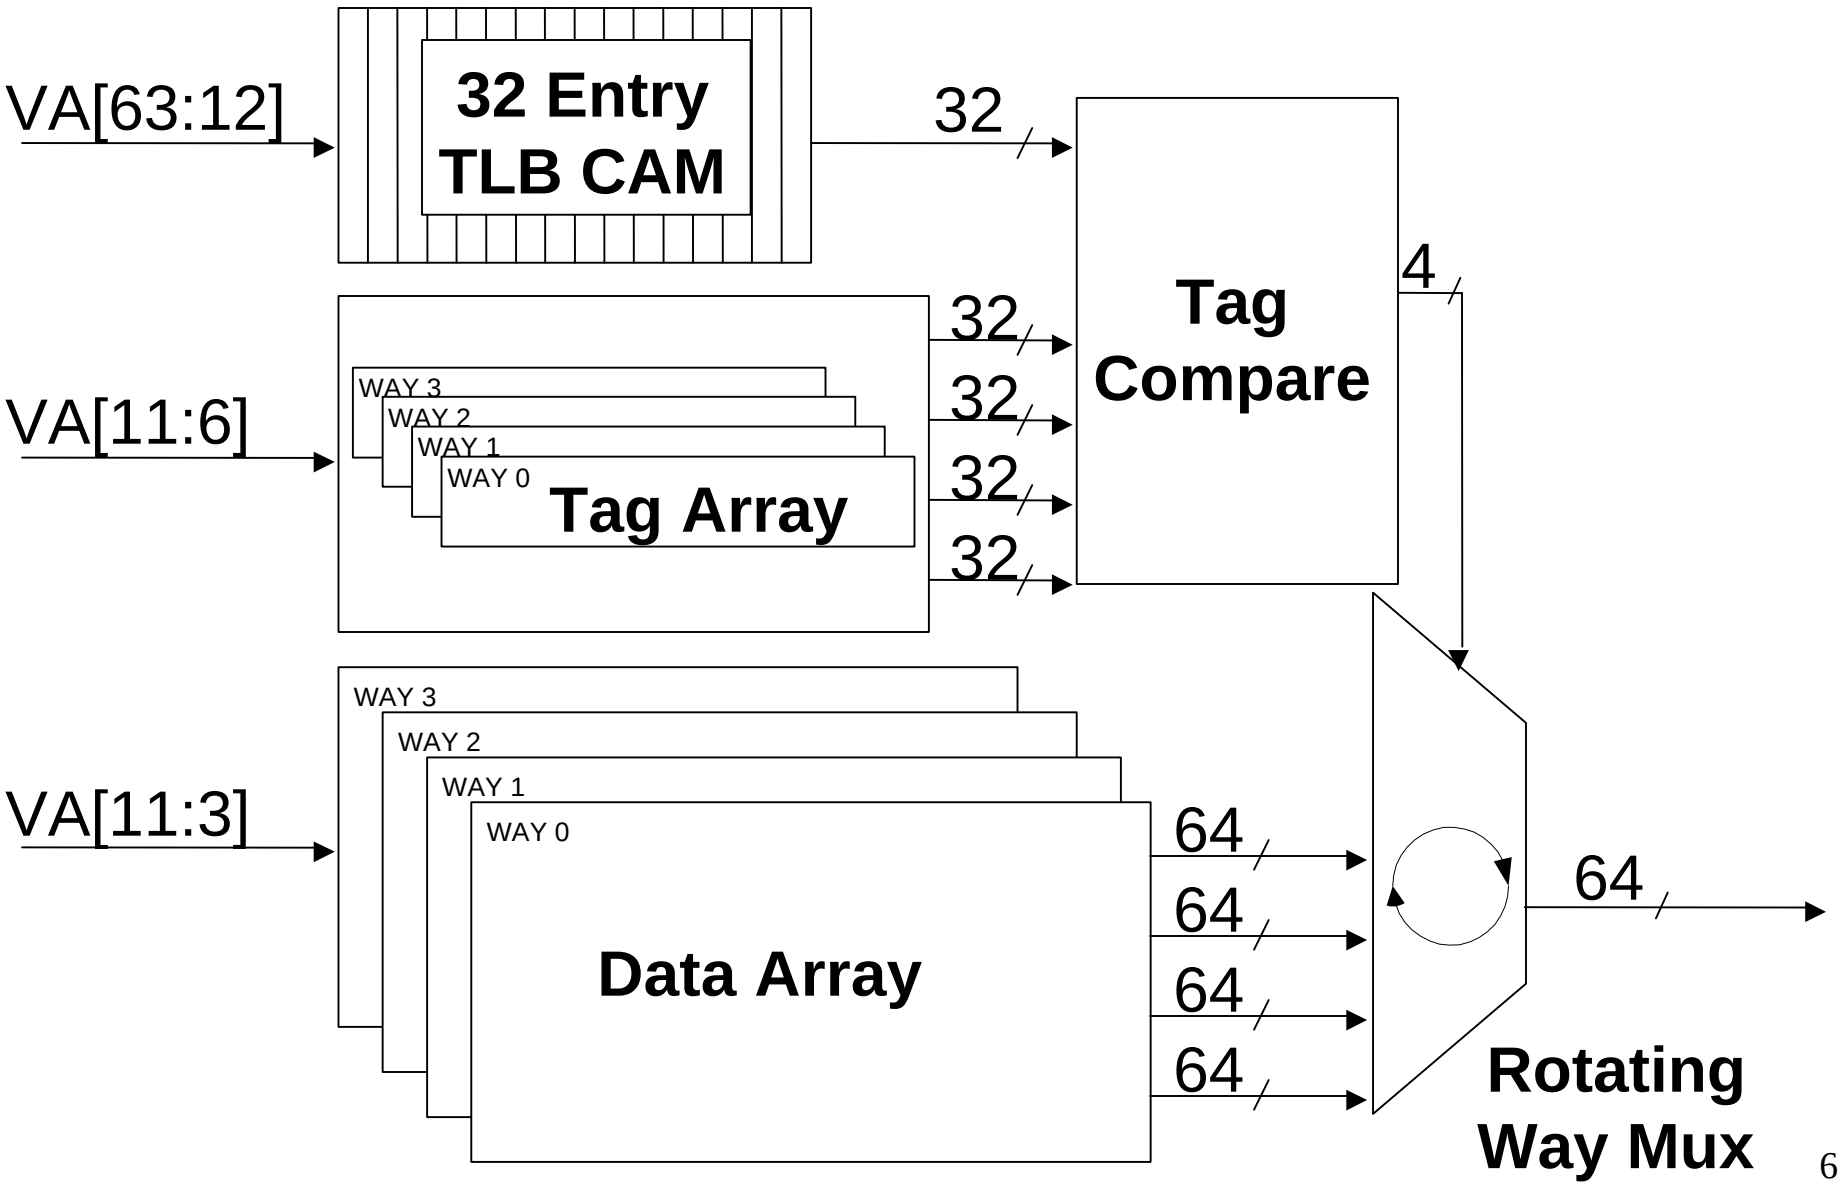

Level 1 Cache Photo

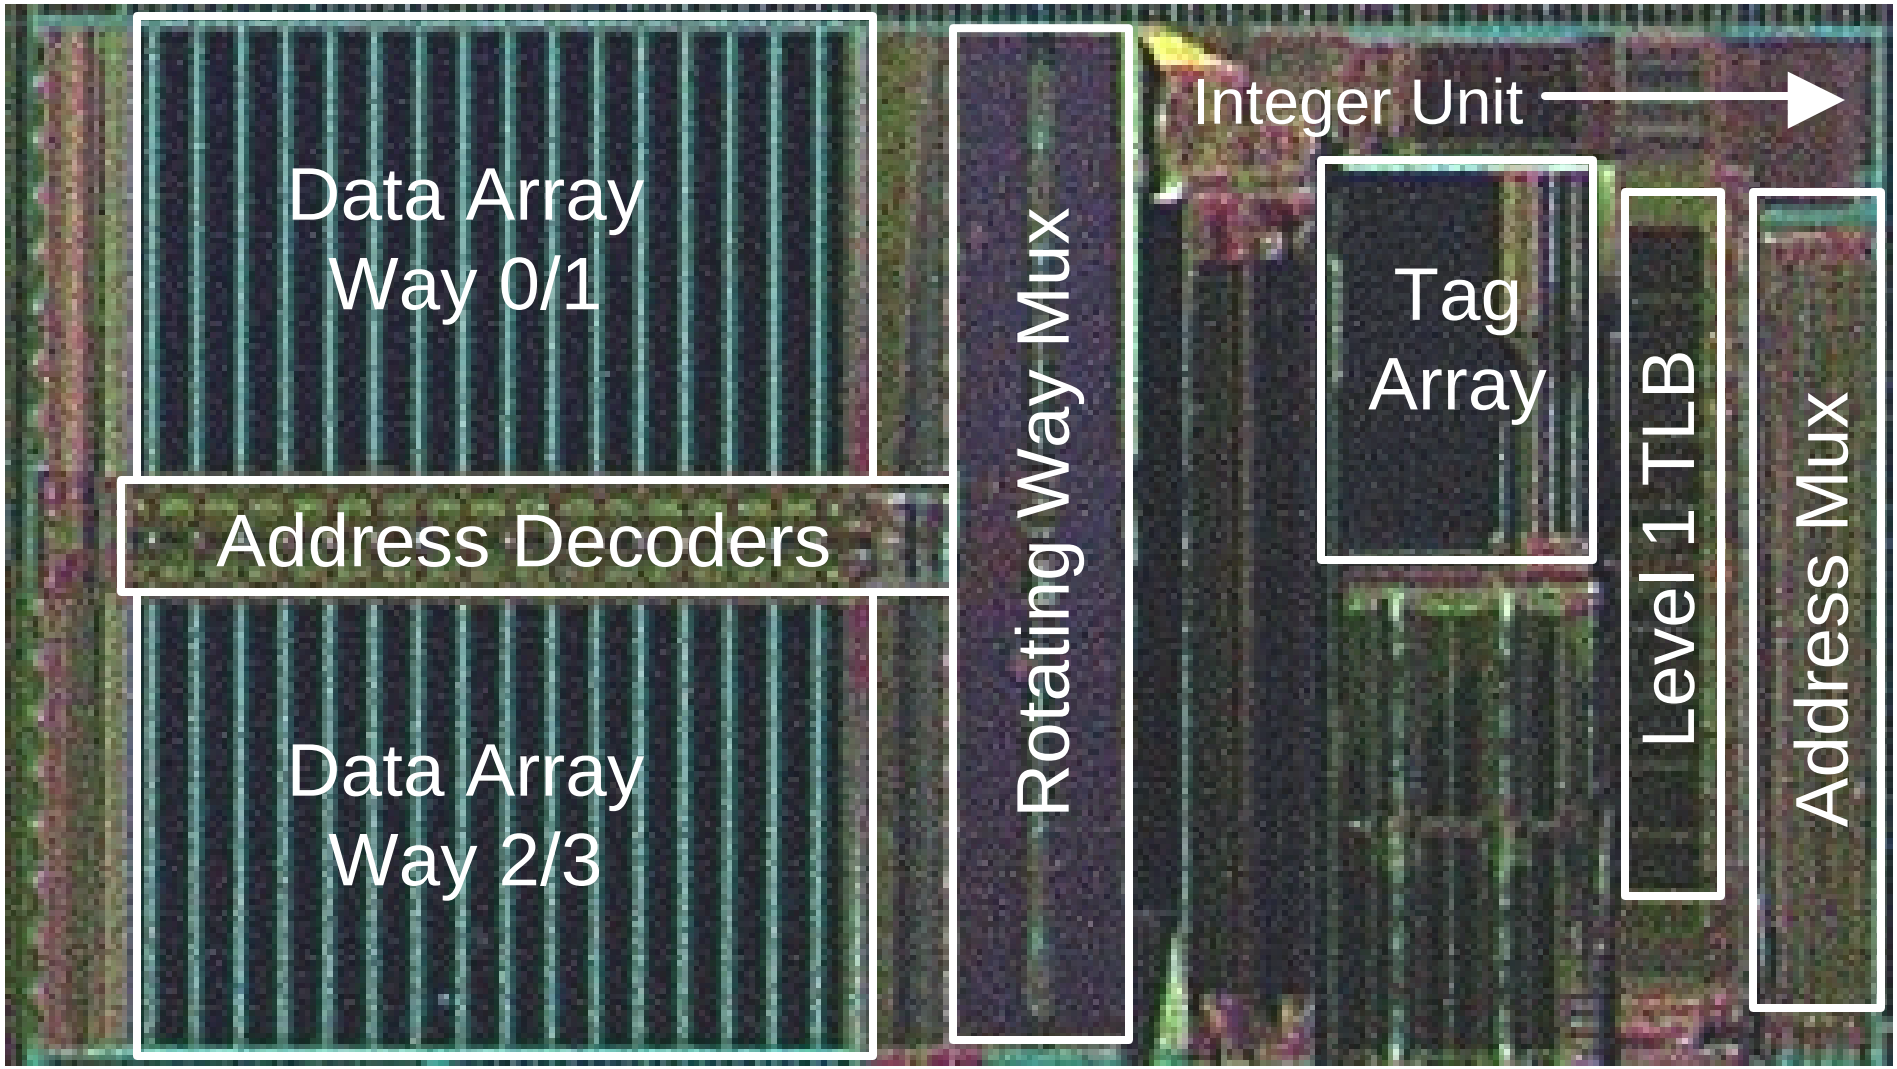

Level 1 TLB Overview

- 32 Entries, each of which stores 115 bits
  - 52 virtual address bits
  - 24 Region ID bits
  - 38 physical address bits
  - 1 Valid bit
- Fixed 4 kB page size
- 2 independent CAM ports, each with shared CAM lines for both virtual and physical address CAMs
  - This saves wires in the critical dimension
- 2-level dynamic CAM with a self timed path to create a monotonic-rising match signal.

Level 1 TLB Pulsed Bit Lines

- The virtual address is presented to the CAM cell on a pair of dual-rail pulsed bit lines.
- Pulsed bit lines eliminate precharge drive fights and allow for the removal of clocked evaluate FETs on downstream dynamic gates.

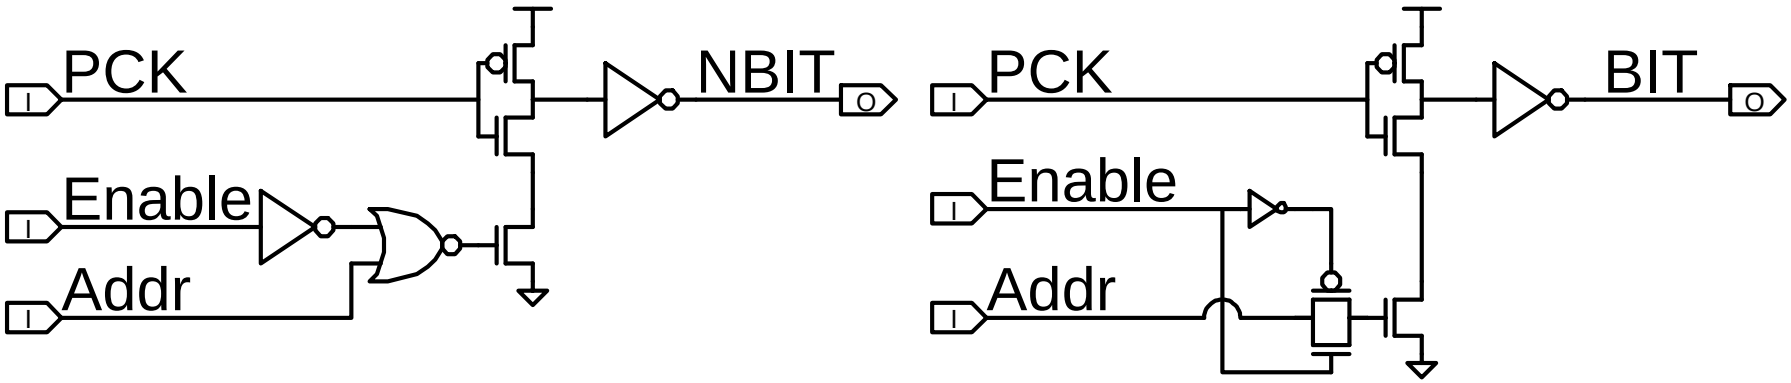

Level 1 TLB CAM Cell

- One write port +2 CAM ports
- CAM bit lines are pulsed
- Cell discharges match[01] when its value does not match the incoming address
- Cell layout is 27.3 μm² (4.8 × standard RAM cell)
- Stability is not an issue, since the pass-fets only turn on during a write.
- Charge sharing is not an issue since 3 of every 4 pulldown stacks don't evaluate.

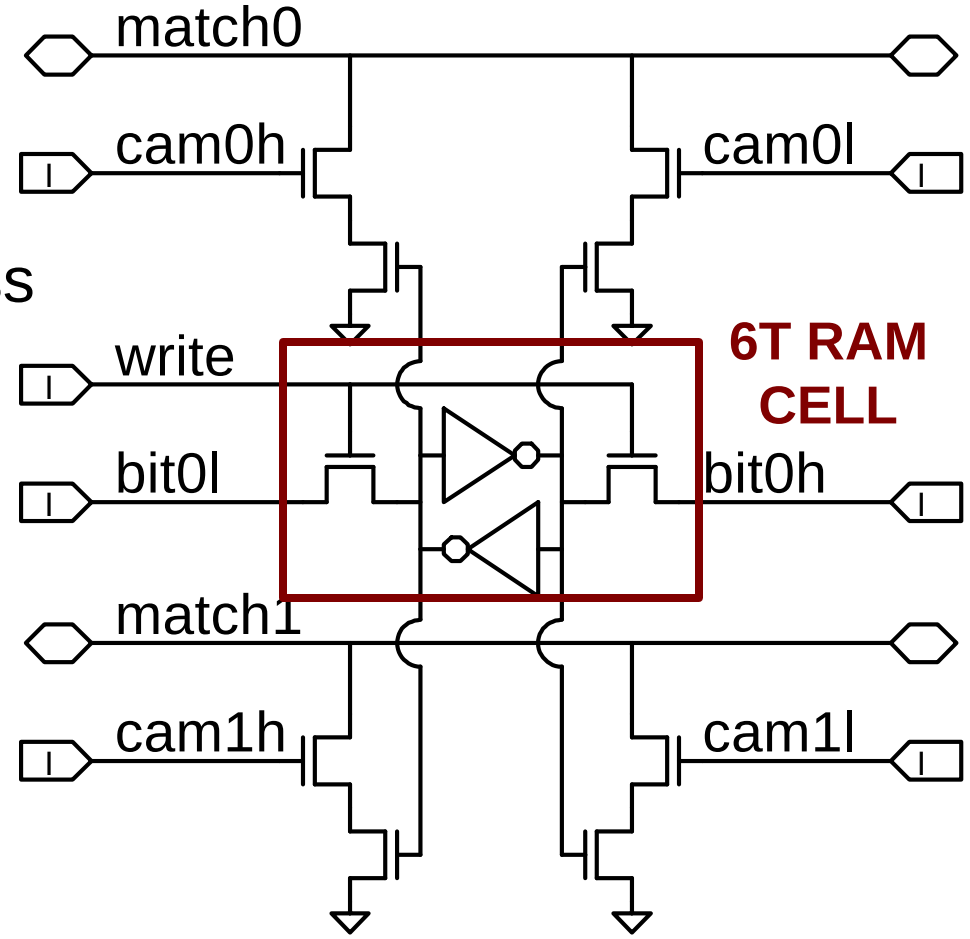

Level 1 TLB Self-Timed Path

- We use a full-rail hierarchical dynamic structure for speed. 6 CAM cells -> 2-input NAND -> 4-wide OR -> 2-input NAND
- This scheme creates a monotonic-falling match signal, but the TAG needs a monotonic-rising match signal. Thus, a “dummy” self-timed path is used to create a hard edge.
- The “Johnson NOR” is used to account for PVT and clock route variations between the two halves of the TLB.
- The self-timed path consumes 20% of the path delay in margin.

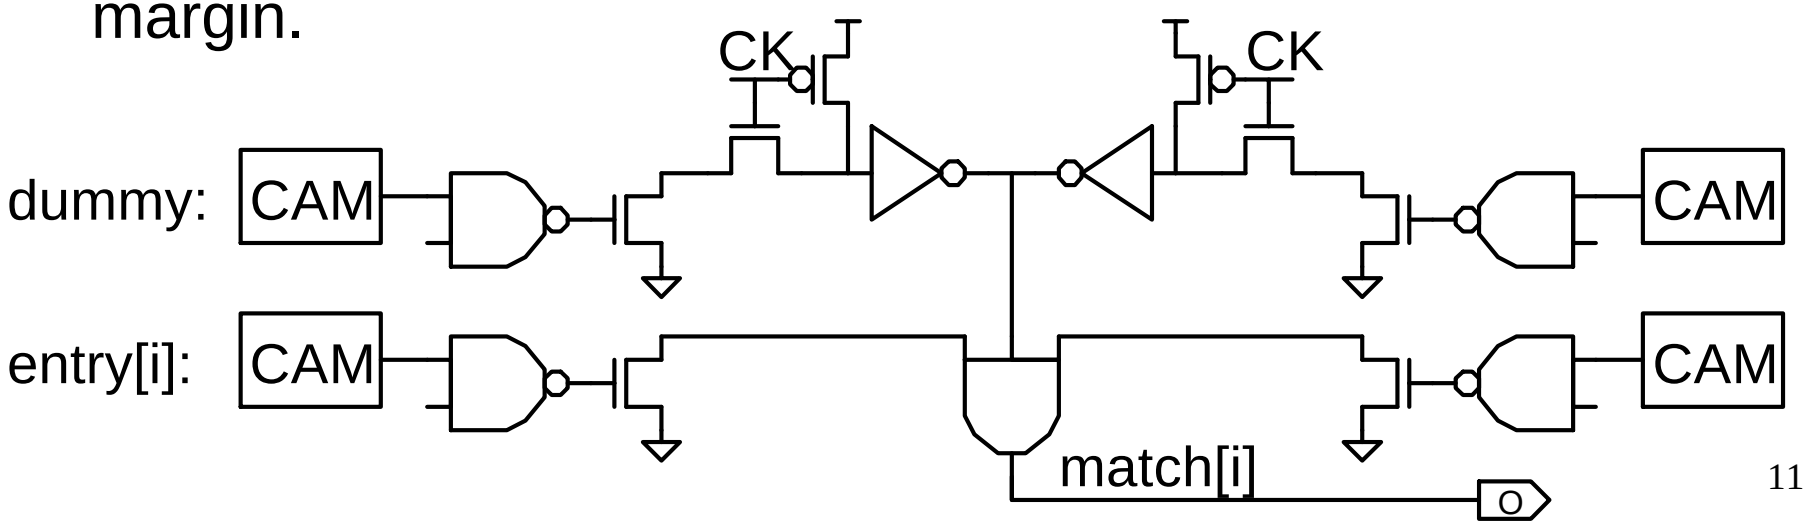

Prevalidated Tag Array

- >1 kB array
  - 64 indices × 4 ways × (32b one-hot vector +1 valid bits)
- 2 independent read ports
- Two-level dynamic full-rail RAM read for speed
  - 8 RAM cells -> inverter -> 8-wide pulldown.
  - No sense-amp.
- SER detection is provided in several ways:
  - Cache is write-through, so spurious invalidations are OK
  - The ways are interleaved so that adjacent multi-cell flips wont appear in the same access.
  - Hardware multi-hot detector for detecting errors.
- 430µm × 700µm in a 0.18µm process.

Prevalidated Tag RAM Cell

- Two read ports + one invalidate port
- Read ports are both on same side of RAM array
  - Generate correct monotonicity for tag compare
- To write the cell, a read port is used to write 1's, while the invalidate port is used to write 0's.
- Cell is asymmetric to speed up the read.
- Cell layout is 11.28 μm^2
  (2.0 × standard RAM cell)

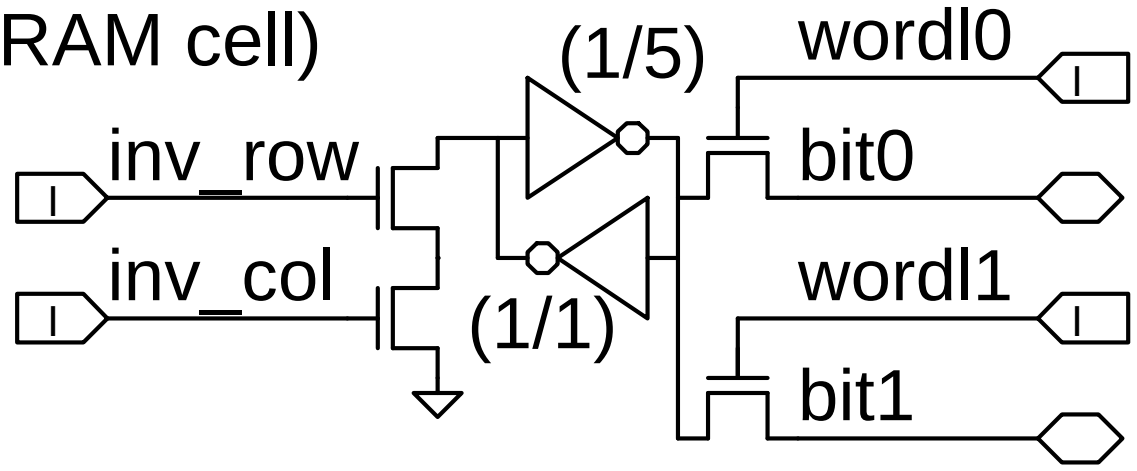

Level 1 Data Array Overview

- 16 kB of data + 2kB of parity
  - 4 ways of 4kB each
- Parity protected at the byte level to allow sub-word stores
- 2 independent load ports and 2 store ports.
- Supports 3 operations (all single cycle)
  - Loads (8B) – Stores (1B - 8B) – Fills (64B)
- Array is approximately 50% RAM-area efficient
  - Tradeoff area for speed.
- 1280µm × 1770µm in a 0.18µm process.

Level 1 Data Array Decoder

- Decode is performed hierarchically.
- Pulsed decoders allow for the removal of clocked-evaluate FETs in downstream dynamic logic.
- Static AND gate provides noise immunity on long pulse routes.

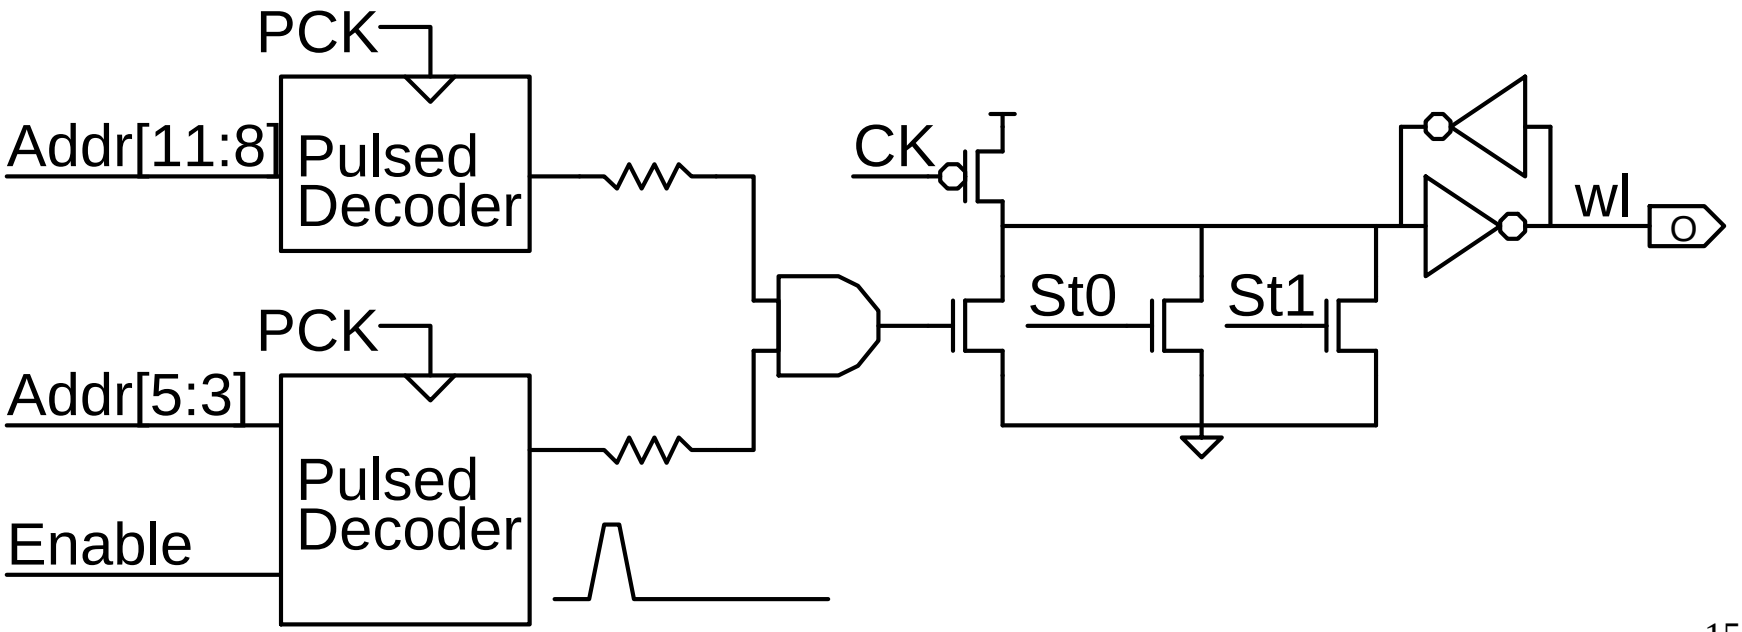

Level 1 Data Array RAM Cell

- 6T-RAM cell is made dual-ported by independently controlling the word lines for the pass fets.
- Single-ended full-rail dynamic reads
- Cell uses large NFETs for speed.
- Stability is afforded by a quick discharge of a lightly loaded bit line during a single ended read.
- Cell layout is 7.01 μm^2 (1.25 × standard RAM cell)

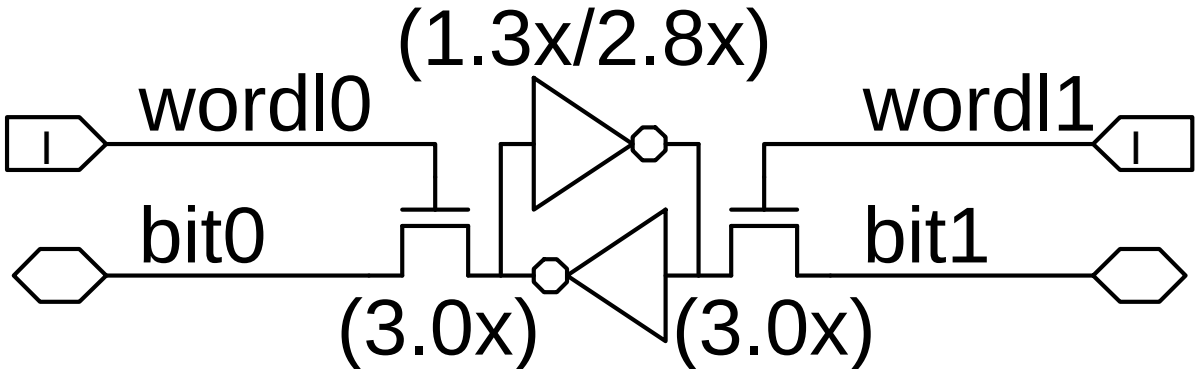

Data Array Local-Global Dump Circuit

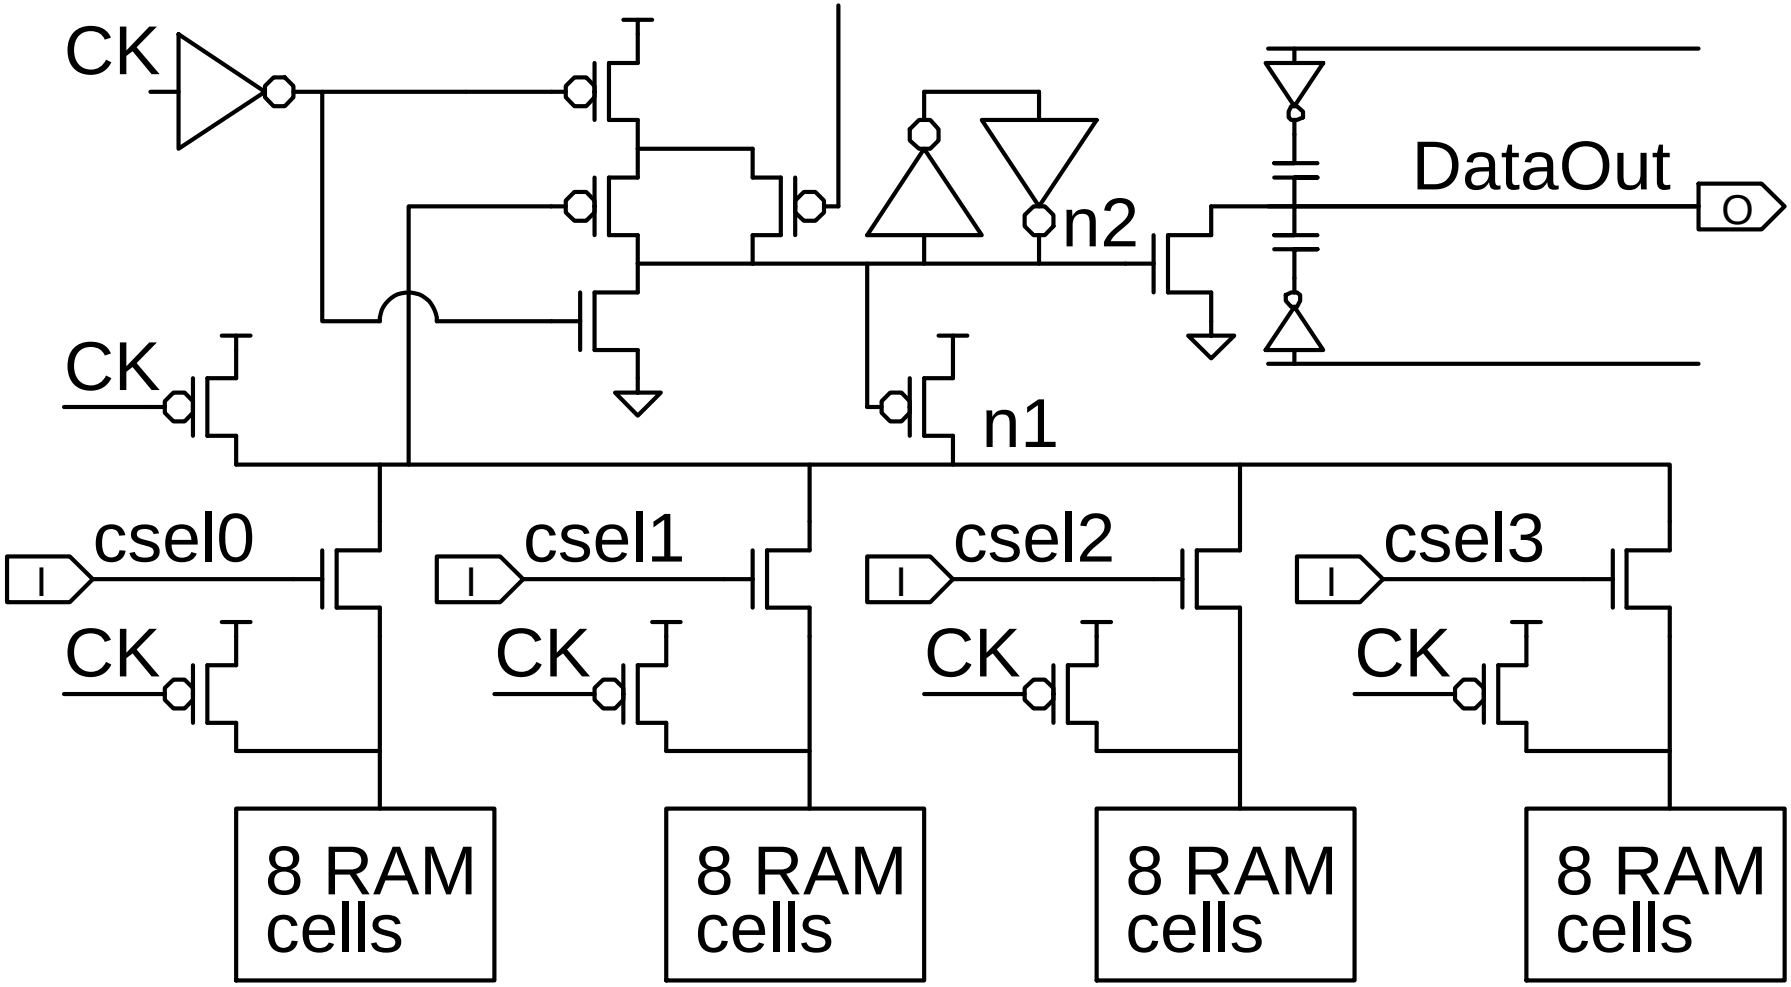

Rotating Way Mux

- This block accomplishes 3 main tasks:
  - Way Selection
  - Byte Rotation (for sub 8B loads)
  - Little Endian / Big Endian Transformations
- A series of monotonicity-preserving static gates are used to perform the first stages of muxing.
  - Provides noise immunity
  - Allows phase stealing
- A pulse-precharged distributed pulldown network is used for the final stage.

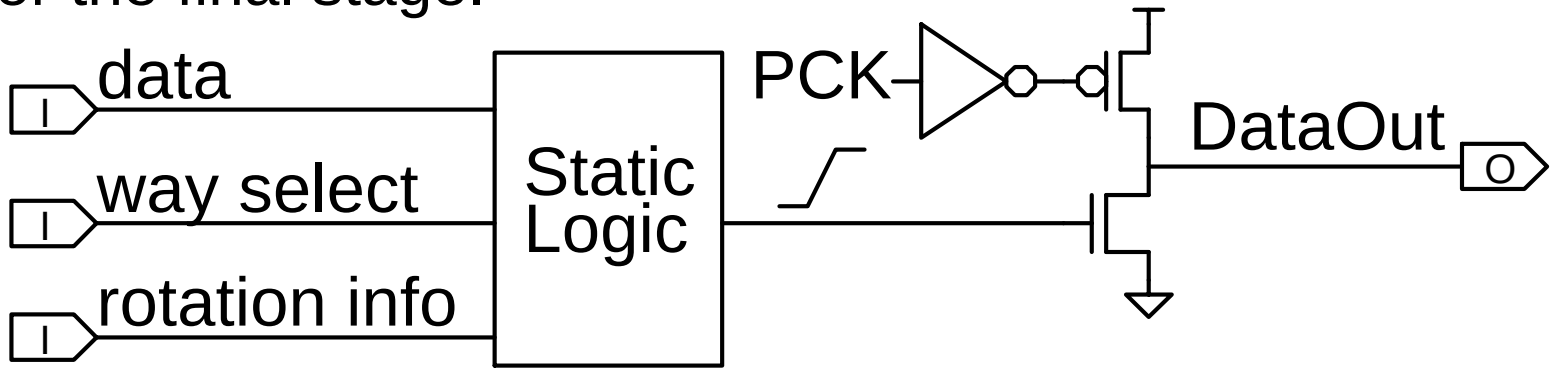

Correlated Performance Data

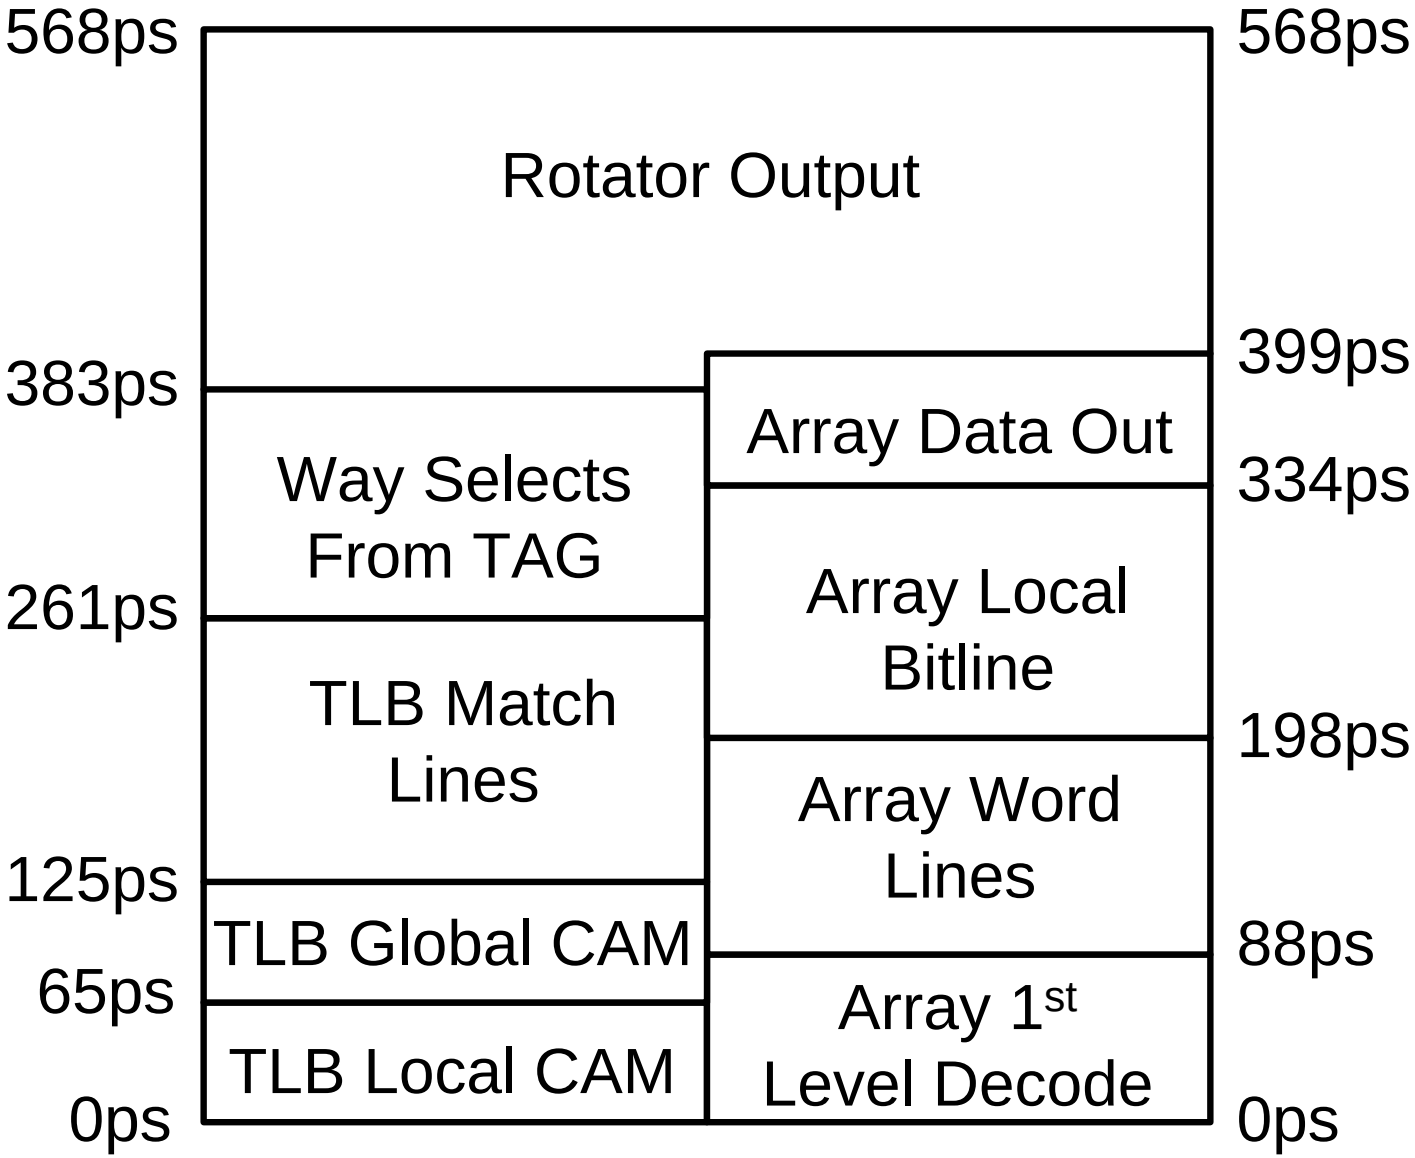

Summary

- Small way size enables TLB, TAG, and data array accesses all to start w/ untranslated address bits.
- Prevalidated tags speed up both the TLB access and the TAG compare.
- Pulsed circuits allow for the removal of clocked evaluate-FETs in downstream dynamic logic.
- Single-rail full-swing hierarchical dynamic RAM reads trade off area for speed.
- All of this combined with a tight layout amounts to a “true” single cycle cache.